

**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES
(Attorney Docket № 13783US02)**

In the Application of:

Uri Elzur et al.

Serial No.: 10/652,330

Filed: August 29, 2003

For: SYSTEM AND METHOD FOR
NETWORK INTERFACING

Examiner: Hoang, Hieu T.

Group Art Unit: 2452

Confirmation No.: 1614

Electronically filed on April 19, 2010

APPEAL BRIEF

Mail Stop Appeal Brief – Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

This is an appeal from an Office Action dated October 6, 2009 ("Final Office Action"), in which claims 18-21, 25 and 33-41 were finally rejected. The Appellant respectfully requests that the Board of Patent Appeals and Interferences ("Board") reverses the final rejection of claims 18-21, 25 and 33-41 of the present application. The Appellant notes that this Appeal Brief is timely filed within the period for reply that ends on April 19, 2010.

REAL PARTY IN INTEREST
(37 C.F.R. § 41.37(c)(1)(i))

Broadcom Corporation, a corporation organized under the laws of the state of California, and having a place of business at 5300 California Avenue, Irvine, California 92617, has acquired the entire right, title and interest in and to the invention, the application, and any and all patents to be obtained therefor, as set forth in the Assignment recorded at Reel 014276 and 014309, Frames 0255 and 0886 in the PTO Assignment Search room.

RELATED APPEALS AND INTERFERENCES
(37 C.F.R. § 41.37(c)(1)(ii))

The Appellant is unaware of any related appeals or interferences.

STATUS OF THE CLAIMS
(37 C.F.R. § 41.37(c)(1)(iii))

The present application includes pending claims 18-21, 25 and 33-41, all of which have been rejected. The Appellant identifies claims 18-21, 25 and 33-41 as the claims that are being appealed. The text of the pending claims is provided in the Claims Appendix.

STATUS OF AMENDMENTS
(37 C.F.R. § 41.37(c)(1)(iv))

The Appellant has not amended any claims subsequent to the final rejection of claims 18-21, 25 and 33-41 mailed on October 6, 2009.

SUMMARY OF CLAIMED SUBJECT MATTER
(37 C.F.R. § 41.37(c)(1)(v))

The Appellant has inserted Figs. 6 and 9 of the present application below, to illustrate several aspects of the invention.

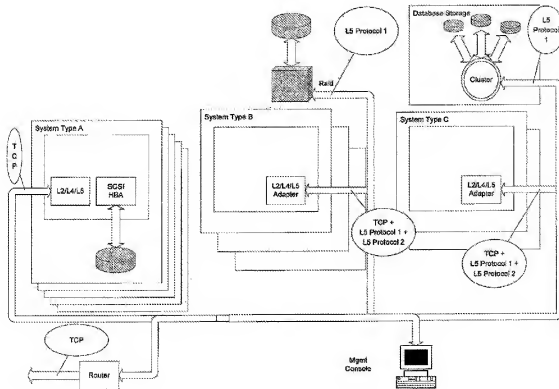


FIG. 6

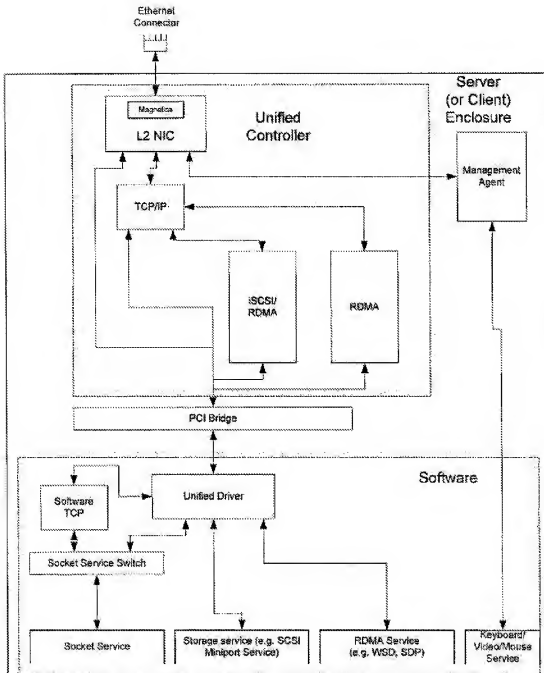


FIG. 9

Independent claim 18 recites the following:

A network communication system¹, comprising:
a single integrated convergent network controller chip²; and
a single Ethernet connector³ for handling a plurality of different types of network traffic⁴ transported via a single fabric⁵, wherein:
the single Ethernet connector is coupled to the single integrated convergent network controller chip⁶;
the single fabric is coupled to a plurality of servers⁷;
the single integrated convergent network controller chip⁸ is operable to concurrently process⁹ the plurality of different types of network traffic¹⁰ for the plurality of servers¹¹, which is transported via the single fabric¹².

¹ See *present specification* at, e.g., Fig. 6, data center.

² See *id.*, e.g., Fig. 6, L2/L4/L5 adapter and Fig. 9, unified controller.

³ See *id.*, e.g., Fig. 9, Ethernet connector.

⁴ See *id.*, e.g., network (TCP/IP) traffic, storage (SAN/NAS) traffic, cluster (IPC) traffic and management (keyboard/video/mouse) traffic; p. 8, ¶39, lines 2-4; Fig. 9, TCP/IP, iSCSI/RDMA, RDMA traffic, management agent traffic.

⁵ See *id.*, e.g., Fig. 6, single fabric coupled to the L2/L4/L5 adapter; p. 8, ¶34, lines 2-4.

⁶ See *id.*, e.g., Fig. 9, Ethernet connector is coupled to the unified controller.

⁷ See *id.*, e.g., Fig. 6, single fabric coupled to system type A-C, RAID storage and cluster data base storage.

⁸ See *id.*, e.g., Fig. 6, L2/L4/L5 adapter and Fig. 9, unified controller.

⁹ See *id.*, e.g., p. 9, ¶38, line 10; p. 14, ¶53, line 1.

¹⁰ See *id.*, e.g., network (TCP/IP) traffic, storage (SAN/NAS) traffic, cluster (IPC) traffic and management (keyboard/video/mouse) traffic; p. 8, ¶39, lines 2-4; Fig. 9, TCP/IP, iSCSI/RDMA, RDMA traffic, management agent traffic.

¹¹ See *id.*, e.g., Fig. 6, single fabric coupled to system type A-C, RAID storage and cluster data base storage.

¹² See *id.*, e.g., Fig. 6, single fabric transporting network (TCP/IP) traffic, storage (SAN/NAS) traffic, cluster (IPC) traffic and management (keyboard/video/mouse) traffic;

Claims 19-21 and 25 are dependent directly or indirectly upon independent claim 18.

Independent claim 33 recites the following:

A method for communication, the method comprising:

in a data center¹³:

accessing a storage system¹⁴ over a single fabric¹⁵, wherein said single fabric comprises a single layer 2 (L2) connector¹⁶ coupled to a single integrated convergent network controller chip¹⁷ that is enabled to concurrently process¹⁸ a plurality of different types of traffic¹⁹; and

accessing one or more of a cluster²⁰ and a network over said single fabric²¹.

Claims 34-35 are dependent directly or indirectly upon independent claim 33.

p. 8, ¶39, lines 2-4; Fig. 9, TCP/IP, iSCSI/RDMA, RDMA traffic, management agent traffic.

¹³ See *id.*, e.g., Fig. 6, data center; p. 9, ¶39, lines 1-2.

¹⁴ See *id.*, e.g., Fig. 6, RAID storage.

¹⁵ See *id.*, e.g., Fig. 6, single fabric coupled to system type A-C, RAID storage and cluster data base storage.

¹⁶ See *id.*, e.g., Fig. 9, Ethernet connector.

¹⁷ See *id.*, e.g., Fig. 6, L2/L4/L5 adapter and Fig. 9, unified controller.

¹⁸ See *id.*, e.g., p. 9, ¶38, line 10; 1p. 14, ¶53, line 1.

¹⁹ See *id.*, e.g., network (TCP/IP) traffic, storage (SAN/NAS) traffic, cluster (IPC) traffic and management (keyboard/video/mouse) traffic; p. 8, ¶39, lines 2-4; Fig. 9, TCP/IP, iSCSI/RDMA, RDMA traffic, management agent traffic.

²⁰ See *id.*, e.g., Fig. 6, cluster data base storage.

²¹ See *id.*, e.g., Fig. 6, single fabric transporting network (TCP/IP) traffic, storage (SAN/NAS) traffic, cluster (IPC) traffic and management (keyboard/video/mouse) traffic; p. 8, ¶39, lines 2-4; Fig. 9, TCP/IP, iSCSI/RDMA, RDMA traffic, management agent traffic.

Independent claim 36 recites the following:

A system²² for communication, the system comprising:

a single integrated convergent network controller chip²³ that enables concurrent²⁴ hardware, firmware and software processing functionalities of a plurality of different types of traffic²⁵ that are received via a single layer 2 (L2) connector²⁶ that is communicatively coupled to a plurality of servers²⁷ via a single fabric²⁸.

Claims 37-38 are dependent directly or indirectly upon independent claim 36.

Independent claim 39 recites the following:

A method for communication, the method comprising:

concurrently providing²⁹, via a single integrated convergent network controller chip³⁰, hardware, firmware and software processing functionalities³¹ of a plurality of

²² See *id.*, e.g., Fig. 6, data center; p. 9, ¶39, lines 1-2.

²³ See *id.*, e.g., Fig. 6, L2/L4/L5 adapter and Fig. 9, unified controller.

²⁴ See *id.*, e.g., p. 9, ¶38, line 10; 1p. 14, ¶53, line 1.

²⁵ See *id.*, e.g., network (TCP/IP) traffic, storage (SAN/NAS) traffic, cluster (IPC) traffic and management (keyboard/video/mouse) traffic; p. 8, ¶39, lines 2-4; Fig. 9, TCP/IP, iSCSI/RDMA, RDMA traffic, management agent traffic.

²⁶ See *id.*, e.g., Fig. 9, Ethernet connector.

²⁷ See *id.*, e.g., Fig. 6, single fabric coupled to system type A-C, RAID storage and cluster data base storage.

²⁸ See *id.*, e.g., Fig. 6, single fabric transporting network (TCP/IP) traffic, storage (SAN/NAS) traffic, cluster (IPC) traffic and management (keyboard/video/mouse) traffic; p. 8, ¶39, lines 2-4; Fig. 9, TCP/IP, iSCSI/RDMA, RDMA traffic, management agent traffic.

²⁹ See *id.*, e.g., p. 9, ¶38, line 10; 1p. 14, ¶53, line 1.

³⁰ See *id.*, e.g., Fig. 6, L2/L4/L5 adapter and Fig. 9, unified controller.

³¹ See *id.*, e.g., p. 10, ¶40, lines 5-6.

different types of traffic³² that are received via a single layer 2 (L2) connector³³ that is communicatively coupled to a plurality of servers³⁴ via a single fabric³⁵.

Claims 40-41 are dependent directly or indirectly upon independent claim 39.

GROUND OF REJECTION TO BE REVIEWED ON APPEAL
(37 C.F.R. § 41.37(c)(1)(vi))

Claims 18-21, 25 and 33-41 are pending in the instant application. Claims 3 and 22-24 have been previously cancelled, and claims 1-2, 4-17, 26-32 and 42-49 are withdrawn due to allegedly being directed to a non-elected invention. Claims 18, 20, 22-29, 31, 32 and 36-41 are rejected under 35 USC 103(a) as unpatentable by USPP 2001/0037406 ("Philbrick"), in view of USPP 2002/0059451 ("Haviv"). Claims 33-35, 30 are rejected under 35 USC 103(a) as being unpatentable over Philbrick in view of Microsoft Winsock Direct and Protocol Offload on SANs, 03/03/2001 ("Microsoft"). See the Final Office Action at pages 2-11. The Appellant identifies claims 18-21, 25 and 33-41 as the claims that are being appealed. The text of the pending claims is provided in the Claims Appendix.

³² See *id.*, e.g., network (TCP/IP) traffic, storage (SAN/NAS) traffic, cluster (IPC) traffic and management (keyboard/video/mouse) traffic; p. 8, ¶39, lines 2-4; Fig. 9, TCP/IP, iSCSI/RDMA, RDMA traffic, management agent traffic.

³³ See *id.*, e.g., Fig. 9, Ethernet connector.

³⁴ See *id.*, e.g., Fig. 6, single fabric coupled to system type A-C, RAID storage and cluster data base storage.

³⁵ See *id.*, e.g., Fig. 6, single fabric transporting network (TCP/IP) traffic, storage (SAN/NAS) traffic, cluster (IPC) traffic and management (keyboard/video/mouse) traffic; p. 8, ¶39, lines 2-4; Fig. 9, TCP/IP, iSCSI/RDMA, RDMA traffic, management agent traffic.

ARGUMENT
(37 C.F.R. § 41.37(c)(1)(vii))

REJECTIONS UNDER 35 U.S.C. § 103

In order for a *prima facie* case of obviousness to be established, the Manual of Patent Examining Procedure, Rev. 6, Sep. 2007 ("MPEP") states the following:

The key to supporting any rejection under 35 U.S.C. 103 is the clear articulation of the reason(s) why the claimed invention would have been obvious. The Supreme Court in *KSR International Co. v. Teleflex Inc.*, 82 USPQ2d 1385, 1396 (2007) noted that the analysis supporting a rejection under 35 U.S.C. 103 should be made explicit. The Federal Circuit has stated that "rejections on obviousness cannot be sustained with mere conclusory statements; instead, there must be some articulated reasoning with some rational underpinning to support the legal conclusion of obviousness."

See the MPEP at § 2142, citing *In re Kahn*, 441 F.3d 977, 988, 78 USPQ2d 1329, 1336 (Fed. Cir. 2006), and *KSR International Co. v. Teleflex Inc.*, 82 USPQ2d at 1396 (quoting Federal Circuit statement with approval). Further, MPEP § 2143.01 states that "the mere fact that references can be combined or modified does not render the resultant combination obvious unless the results would have been predictable to one of ordinary skill in the art" (citing *KSR International Co. v. Teleflex Inc.*, 82 USPQ2d 1385, 1396 (2007)). Additionally, if a *prima facie* case of obviousness is not established, the Appellant is under no obligation to submit evidence of nonobviousness:

The examiner bears the initial burden of factually supporting any *prima facie* conclusion of obviousness. If the examiner does not produce a *prima facie* case, the applicant is under no obligation to submit evidence of nonobviousness.

See MPEP at § 2142.

Obviousness also requires that the Examiner provide “some articulated reasoning with some rationale underpinning to support the legal conclusion of obviousness.” See *KSR International Co. v. Teleflex Inc.*, 127 S. Ct. 1727, 1741 (2007) quoting *In re Kahn*, 441 F.2d 997,988 (CA Fed. 2006). Put another way, the Examiner should “identify a reason that would have prompted a person of ordinary skill in the relevant field to combine the elements in the way the claimed new invention does.” *KSR*, 127 S. Ct. at 1741. The Examiner should make “explicit” this rationale of “the apparent reason to combine the known elements in the fashion claimed,” including a detailed explanation of “the effects of demands known to the design community or present in the marketplace” and “the background knowledge possessed by a person having ordinary skill in the art.” *Id.*

I. The Proposed Combination of Philbrick and Haviv Does Not Render Claims 18, 20-21, 25 and 36-41 Unpatentable

The Appellant turns to the rejection of claims 18, 20-21, 25 and 36-41 under 35 U.S.C. § 103(a) as being unpatentable over Philbrick in view of Haviv.

A. Independent Claims 18, 36 and 39

With regard to the rejection of independent claim 18 under 35 U.S.C. § 103(a), the Appellant submits that the combination of Philbrick and Haviv does not disclose or suggest “a single Ethernet connector for handling a plurality of different types of

network traffic transported via a single fabric, wherein the single Ethernet connector is coupled to the single integrated convergent network controller chip," or "... the single integrated convergent network controller chip is operable to concurrently process the plurality of different types of network traffic for the plurality of servers, which is transported via the single fabric" as recited in Appellant's claim 18. In order to clarify the argument, the Appellant has inserted Philbrick's Figs. 6 and 14-16 for reference:

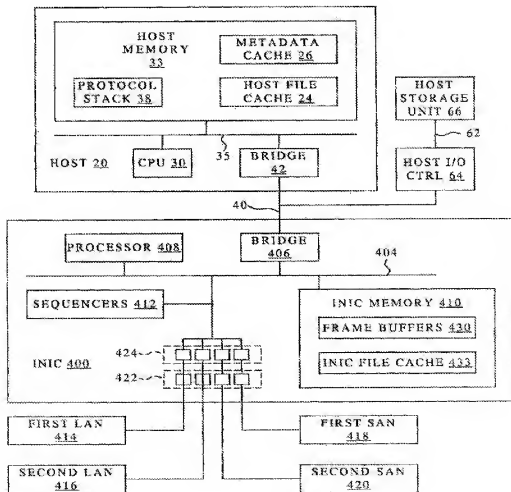


FIG. 6




FIG. 14

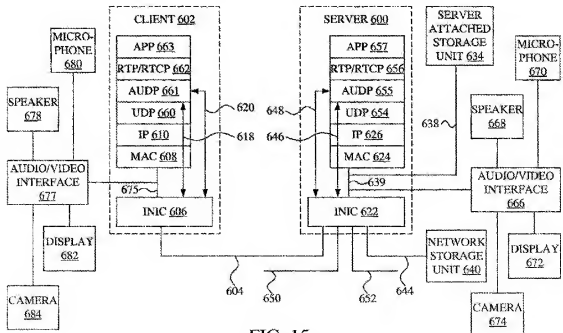


FIG. 15

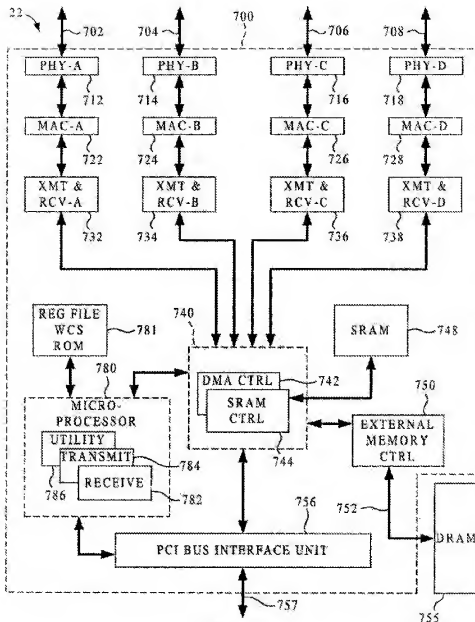


FIG. 16

In the Final Office Action, the Examiner states the following:

"For claim 18, Philbrick discloses a server, comprising: a **single integrated convergence network controller chip** (fig. 6, fig. 1, **network interface card INIC 22**); in a single Ethernet connector for handling a plurality of different types of network traffic (fig. 16, one of the Ethernet

connectors for receiving multiple traffic types, [0065], SCSI and TCP, or Ether storage or SEP and TCP, [0069] lines 20-23, ...”

See Final Office Action at page 5 (emphasis added). The Appellant respectfully points out that Philbrick (see Philbrick's Fig. 6, 16 and at ¶10106) clearly discloses that the INIC 400 (the alleged “single integrated convergent network controller chip”) is connected to **four network lines 702-708**, via the MAC 424 to **four corresponding network connectors PHY/MAC 712-718/722-728**. In this regard, Philbrick does not disclose that an alleged single connector (i.e., a single PHY/MAC) is coupled to the alleged single integrated convergent network controller chip (i.e., INIC 22) via an alleged single fabric (i.e., a single network line).

In addition, the Appellant maintains the arguments at page 21 of the 12/12/08 RCE response, that an “Ethernet connector” is merely a passive mechanical hardware for interfacing IEEE 802.x specifications network signals. The claimed “Ethernet connector” does not utilize active electronic hardware such as ICs and program code to perform signal processing functions, which Philbrick's MAC controller 424 does. In this regard, the Examiner is erroneous in equating Philbrick's MAC controller 424 to the claimed “Ethernet connector”, as recited in claim 18.

The Examiner (see Final Office Action in page 2) argued that Appellant's claim does not exclude using more than one Ethernet connector. The Appellant disagrees, since Appellant's claim 18 clearly recites using only a “**single Ethernet connector for handling a plurality of different types of network traffic transported via a single fabric**”.

In addition, the Examiner (*see* Final Office Action in page 2) also relied for support on Philbrick's Fig. 14 and argued that there is no need for Philbrick to have four connectors to function, and that Philbrick discloses at least two different types of traffic (e.g., iSCSI and TCP/NetBios/SMB) over the same fabric. The Appellant respectfully disagrees with the Examiner's allegation. If assuming the Examiner's above allegation is true, then Philbrick's Fig. 14 should at least disclose that the host server 600 is connected to the client (server) 602, the NAS storage unit 642 and the SAN storage unit 640 **via a single network line 644 (i.e., via a single Ethernet connector and a single fabric)**. However, Philbrick does not make such disclosure.

Contrary to the Examiner's allegation, Philbrick's Fig. 14 clearly discloses that the host server 600 is using separate network lines 604 and 644 (i.e., more than a single Ethernet connector and fabric) to communicate with the client (server) 602, the NAS storage unit 642 and SAN storage unit 640, respectively. In this regard, the Examiner's allegation that "there is no need for Philbrick to have four (i.e., more than a single) connectors to function..." is unsupported by the Philbrick reference. Therefore, the Appellant maintains that Philbrick does not disclose or suggest "a single Ethernet connector for handling a plurality of different types of network traffic transported via a single fabric, wherein the **single** Ethernet connector is coupled to the single integrated convergent network controller chip ...," as recited in Appellant's claim 18. Haviv does not overcome Philbrick's above deficiencies.

In the Final Office Action, the Examiner also states the following:

"...different storage protocols over TCP/IP, [0084], [0085], NAS traffic and network storage traffic over network line 644, utilizing iSCSI and TCP/NetBios/SMB, [0085], iSCSI and TCP/NetBios/SMB, fig. 15, [0093], [0097], [0099], fast path audio and video traffics and real time voice/video traffics and NAS, RTP/RTCP and SIP and MGCP)), the single Ethernet connector is coupled to the single integrated convergent network controller chip ([0066] lines 12-15, Ethernet connector 424 coupled to the INIC)"

See Final Office Action at page 5 (emphasis added). The Examiner seemed to allege that Philbrick's Fig. 14 and ¶¶0084-0085 disclose that the INIC 622 (the alleged "single integrated convergent network controller chip") handles at least two types of traffic over the same fabric (i.e., network line 644), namely the iSCSI SAN traffic and TCP/NetBios/SMB NAS traffic. The Appellant respectfully disagrees, and points out that even though Philbrick discloses that the same network line 644 (the alleged "single fabric") is connected to both the SAN storage unit 640 and the NAS storage unit 642, Philbrick nevertheless, discloses that **only one storage unit is accessed at a time, not concurrently**. In other words, there is still only one type of traffic, **either** the iSCSI SAN traffic **or** the TCP/NetBios/SMB NAS traffic, but **not both** are handled by the network line 644 (the alleged "single fabric"). The Examiner is referred to the following citation of Philbrick:

"A storage fast-path is provided by the INIC 622, under control of the server, for data transferred between network storage units 640 or 642 and client 602 that does not cross the I/O bus. Data is communicated between INIC 622 and network storage unit 640 in accordance with a block format, such as SCSI/TCP or iSCSI, whereas data is communicated between INIC 622 and NAS storage unit 642 in accordance with a file format, such as TCP/NetBios/SMB. For either storage fast-path the INIC 622 may hold another CCB defining a connection with storage unit 640 or 642."

See Philbrick at ¶0085 (emphasis added). As seen in the above citation, Philbrick clearly discloses that **either** the iSCSI data or the TCP/NetBios/SMB data is transferred to the INIC 622. In this regard, **the network line 644** (the alleged "single fabric") still **handles one type of traffic at a time, but not concurrently**.

In addition, the Examiner had argued before, that the iSCSI traffic includes two separate protocol types, namely, the SCSI and TCP network protocols, and therefore Philbrick's INIC allegedly performs "**concurrently processing the plurality of different types of network traffic**". The Appellant respectfully disagrees, and refers the Examiner to Appellant's arguments in the 7/23/09 response (see pages 16-17), that the iSCSI traffic is a **single traffic type**, not two traffic types. Specifically, the Examiner is referred to Satran, which discloses that the **SCSI** (not **iSCSI**) protocol and commands by themselves cannot be transported across the TCP/IP network infrastructure. In this regard, Satran clearly discloses that the **iSCSI protocol is a stand alone new protocol standard, with its own commands and numbering scheme for the iSCSI PDUs, which are transported over the TCP/IP network infrastructure**.

Therefore, contrary to the Examiner's allegation, the iSCSI protocol by itself, is a unique protocol type (i.e., a single protocol of a single traffic type), and is not considered as two separate protocol types (i.e., separate SCSI and TCP network traffics). In this regard, Philbrick's INIC handles the iSCSI traffic as a **single traffic type only**, and not two traffic types, as alleged by the Examiner. Accordingly, Philbrick's INIC does not

disclose or suggest “concurrently processing the plurality of different types of network traffic,” as recited in Appellant’s claim 18.

The Examiner further argued the following:

“In response to arguments of Philbrick’s INIC does not process a plurality of traffics concurrently. In traversal, Philbrick clearly discloses that the **server INIC holds 3 CCB’s concurrently for distinguishing 3 different types of network traffics**, therefore supporting a chip concurrently process a plurality of traffic types ([0085], 3 CCBs for different traffic types so that the server INIC can process the traffic types according to the server protocol stack in fig. 14).”

See Final Office Action in page 4 (emphasis added). The Examiner alleged that Philbrick’s Fig. 14 and ¶0085 disclose that the INIC 622 allegedly **concurrently** holds three CCBs (i.e., **client CCB**, **SAN CCB** and **NAS CCB**) for distinguishing three different types of network traffic (i.e., TCP, iSCSI and SMB), therefore allegedly supporting a chip that concurrently processes a plurality of traffic types. **The Appellant points out that the Examiner seems to have misinterpreted Philbrick’s ¶0085, consequently, the Examiner’s above allegation is contrary to Philbrick’s disclosure.** The Examiner is respectfully referred the above citation of Philbrick in ¶0085:

“A storage fast-path is provided by the INIC 622, under control of the server, for **data transferred between network storage units 640 or 642** ... For either storage fast-path the INIC 622 may hold another CCB defining a connection with storage unit 640 or 642.”

As seen, Philbrick clearly disclose that the INIC 622 controls data transfer between **SAN** network storage unit 640 or the **NAS** storage unit 642. In addition,

Philbrick also clearly discloses that the INIC 622 holds only one CCB connection at a time, namely, a NAS CCB for the NAS storage unit 640, or a SAN CCB for the SAN storage unit 642.

In this regard, contrary to the Examiner's allegation, **Philbrick does not disclose or suggest that the SAN CCB and the NAS CCB are concurrently handled by the INIC 622.** Moreover, the Appellant also points out that Philbrick's **client CCB is handled on a separate network line** (not in the same alleged "single fabric") and using a separate connector. Therefore, the client CCB is not even handled by the same alleged "single connector" or the alleged "single fabric". Accordingly, Philbrick's INIC 622 does not process all three CCBs (i.e., client CCB, SAN CCB and NAS CCB) concurrently (in a single fabric and using a single Ethernet connector).

Based on the foregoing rationale, the Appellant maintains that Philbrick does not disclose or suggest **"... the single integrated convergent network controller chip is operable to concurrently process the plurality of different types of network traffic for the plurality of servers, which is transported via the single fabric,"** as recited in Appellant's claim 18. Haviv does not overcome Philbrick's above deficiencies.

The Examiner in the Final Office Action (*see* page 6) further alleges that Philbrick (*see* Fig. 15 and ¶¶0093, 0097, 0099) discloses other traffics, such as the fast path audio and video traffics, as well as real time voice/video and NAS RTP/RTCP, SIP and MGCP that are communicated to the client server 602 **via a single fabric.**

The Appellant respectfully disagrees, and points out that the listed traffics (i.e., fast path audio and video traffics, as well as real time voice/video and NAS RTP/RTCP, SIP and MGCP) are **not** communicated to the client server 602 via an alleged "**fabric**". Instead, Philbrick (*see* Fig. 15 and ¶¶0093, 0097, 0099) discloses the traffic is communicated via **an I/O bus 675**, which is neither coupled to an Ethernet connector nor to the INIC 606 of the client server 602.

In this regard, Philbrick's (*see* Fig. 15 and ¶¶0093, 0097, 0099) traffic in the I/O bus 675 does not read on "a single Ethernet connector for handling a plurality of different types of network traffic transported via a single fabric, wherein the single Ethernet connector is coupled to the single integrated convergent network controller chip ...," as recited in Appellant's claim 18. Haviv does not overcome Philbrick's above deficiencies.

Based on the foregoing rationale, the Appellant maintains that Philbrick does not disclose or suggest "the single integrated convergent network controller chip is operable to **concurrently process** the plurality of different types of network traffic (via the single fabric and a single Ethernet connector) **for the plurality of servers**," or "... **the single integrated convergent network controller chip is operable to concurrently process the plurality of different types of network traffic** for the plurality of servers, which is transported via the single fabric," as recited in Appellant's claim 18.

Even though Haviv discloses an interconnected fabric (the alleged “single fabric”), Haviv, nevertheless, still does not overcome Philbrick’s above deficiencies. In this regard, Appellant’s claim 18 is submitted to be allowable. Independent claims 26, 36 and 39 are similar in many respects to the method disclosed in independent claim 18, are also allowable for the reasons stated above with regard to claim 18.

B. Dependent Claims 20-21, 25, 37-38 and 40-41

Based on at least the foregoing, the Appellant believes the rejection of independent claims 18, 36 and 39 under 35 U.S.C. § 103(a) as being unpatentable by the combination of Philbrick and Haviv has been overcome and requests that the rejection be withdrawn. Additionally, claims 20-21, 25, 37-38 and 40-41 depend directly or indirectly from independent claims 18, 36 and 39 and are, consequently, also respectfully submitted to be allowable.

B(1). Rejection of Dependent Claims 25, 38 and 41

The Examiner states the following in the Final Office Action:

“For claim 25, Philbrick-Haviv further discloses the plurality of different types of traffic comprises at least two of network traffic, storage traffic, interprocess communication (IPC) traffic and cluster traffic (Philbrick, [0065] lines 15-21, network traffic TCP/IP and storage traffic SCSI, Haviv, [0019]).”

See Final Office Action in page 7. The Examiner is referred to Appellant’s above arguments in claim 18, namely, that the iSCSI traffic is a single type of traffic

comprising iSCSI PDU messages. Also, Philbrick's Fig. 14 and ¶0085 disclose that only one type of network traffic, **either the SAN or NAS traffic** is handled by NIC 622 (the alleged "integrated chip"). In this regard, Philbrick does not disclose that "at least two of network traffic, storage traffic, IPC traffic or the cluster traffic", are **concurrently** handled by the Ethernet connector and the integrated chip. Claim 25 is therefore allowable. Claims 38 and 41 are allowable for the same rationale as stated with regard to claim 25.

B(2). Rejection of Dependent Claims 37 and 40

The Examiner states the following in the Final Office Action:

"For claim 37, Philbrick-Haviv discloses said single integrated convergent network controller chip comprises a layer 2 network interface card (L2 NIC) (Philbrick, [0065] lines 7-11, Ethernet, fig. 24, MAC controller), a transmission control protocol (TCP) processor, an iSCSI processor ([0065] lines 15-21, iSCSI processing over TCP/IP) and a remote direct memory access (RDMA) processor (fig. 25, DMA controller), and a Management Agent processor ([0106], last sentence)."

See Final Office Action in page 8. The Examiner relies for support on Philbrick's Fig. 6 and [0065] and alleges that Philbrick's INIC 400 discloses a TCP processor, an iSCSI processor. The Appellant respectfully disagrees, and refers the Examiner to the following citation of Philbrick:

"SANS 418 and 420 may run a storage protocol such as SCSI over TCP/IP or SCSI Encapsulation Protocol. One such storage protocol is ... "iSCSI (Internet SCSI) June 2000, which in an earlier Internet-Draft was entitled "SCSI/TCP (SCSI over TCP)," . employs SCSI Encapsulation Protocol (SEP) at the session layer, and either TCP or SAN transport

protocol (STP) at the transport layer, depending primarily upon whether data is being transferred over a WAN or the Internet, for which TCP is used, or **data is being transferred over a LAN or SAN**, for which STP is used."

See Philbrick in ¶0065 (emphasis added). The Appellant points out that Philbrick, in the above citation, discloses that it is the SANS 418 and 420, not the NIC 400 in the host server, which run the iSCSI protocol (i.e., SCSI over TCP). In this regard, Philbrick's Fig. 6 and ¶0065 does not disclose a TCP processor or iSCSI processor. On the contrary, Philbrick's Fig. 14 discloses that the TCP protocol stack is processed in the host, and the SCSI protocol is processed in the storage unit 634.

In addition, Philbrick's Fig. 25 merely discloses a DMA controller 2206, which performs DMA operation. However, Philbrick's Fig. 25, nor in the entire reference, disclose RDMA protocol processing. In this regard, Philbrick does not disclose or suggest the alleged "RDMA protocol processor" in the INIC, as alleged by the Examiner.

Likewise, the Examiner alleges that Philbrick's ¶0106 discloses that the INIC includes a management agent processor. The Appellant respectfully disagrees, and points out that Philbrick's ¶0106 states:

"The MAC units 722, 724, 726 and 728 also provide statistical information that can be used for simple network management protocol (SNMP)."

Philbrick merely discloses that the MAC units **provide statistical information that can be used for SNMP**. There is no disclosure or suggestion that the SNMP protocol is actually processed within the INIC. In this regard, the Examiner's allegation that Philbrick discloses a management agent processor in the INIC, is unsupported.

Based on the foregoing rationale, the Appellant maintains that Philbrick at least does not disclose or suggest “wherein said single integrated convergent network controller chip comprises a layer 2 network interface card (L2 NIC), **a transmission control protocol (TCP) processor, an iSCSI processor, a remote direct memory access (RDMA) processor and a Management Agent processor,**” as recited in claim 37. Haviv does not overcome Philbrick’s above deficiencies. Claim 37 is therefore allowable. Claim 40 is allowable for the same rationale as stated with regard to claim 37.

II. The Proposed Combination over Philbrick and Microsoft Does Not Render Claims 33-35 Unpatentable

The Appellant turns to the rejection of claims 33-35 under 35 U.S.C. § 103(a) as being unpatentable over Philbrick and further in view of Microsoft.

A. Independent Claim 33

With regard to the rejection of independent claim 33 under 35 U.S.C. § 103(a), the Appellant refers the Examiner to the arguments in claim 1, namely, A) Philbrick does not disclose or suggest “a single integrated convergent network controller chip that is enabled to **concurrently process a plurality of different types of traffic**”; and B) Philbrick does not disclose or suggest “**a single connector coupled to a single integrated convergent chip** that is enabled to concurrently process a plurality of different types of traffic”.

Microsoft does not overcome the deficiencies of Philbrick. Accordingly, a prima facie case of obviousness cannot be established by the combination of Philbrick and Microsoft to reject claim 33, therefore, claim 33 should be allowable. The Appellant respectfully requests that the rejection of claim 33 under 35 U.S.C. § 103(a) be withdrawn.

B. Rejection of Dependent Claims 34-35

Based on at least the foregoing, the Appellant believes the rejection of the independent claim 33 has been overcome. Additionally, claims 34-35 depend from independent claim 33, and are, consequently, also respectfully submitted to be allowable.

The Appellant reserves the right to argue additional reasons beyond those set forth herein to support the allowability of dependent claims 18-21, 25 and 33-41, should such a need arise.

CONCLUSION

For at least the foregoing reasons, the Appellant submits that claims 18-21, 25 and 33-41 are in condition for allowance. Reversal of the Examiner's rejection and issuance of a patent on the application are therefore requested.

The Commissioner is hereby authorized to charge \$540 (to cover the Brief on Appeal Fee) and any additional fees or credit any overpayment to the deposit account of McAndrews, Held & Malloy, Ltd., Account No. 13-0017.

Respectfully submitted,

Date: April 19, 2010

/ Frankie W. Wong /

Frankie W. Wong
Registration No. 61,832
Patent Agent for Appellant

McANDREWS, HELD & MALLOY, LTD.
500 WEST MADISON STREET, 34TH FLOOR
CHICAGO, ILLINOIS 60661
(312) 775-8093 (FWW)
Facsimile: (312) 775 – 8100

CLAIMS APPENDIX
(37 C.F.R. § 41.37(c)(1)(viii))

18. A network communication system, comprising:
a single integrated convergent network controller chip; and
a single Ethernet connector for handling a plurality of different types of network traffic transported via a single fabric, wherein:

the single Ethernet connector is coupled to the single integrated convergent network controller chip;

the single fabric is coupled to a plurality of servers;

the single integrated convergent network controller chip is operable to concurrently process the plurality of different types of network traffic for the plurality of servers, which is transported via the single fabric.

19. The network communication system according to claim 18, wherein:
the plurality of servers comprises a blade server, and
the single integrated convergent network controller chip is part of a blade mounted in the blade server.

20. The network communication system according to claim 19, wherein the blade server has a single Internet protocol (IP) address.

21. The network communication system according to claim 18, wherein the plurality of servers is part of a data center, and the data center comprises a plurality of other servers coupled to each other via the single fabric.

25. The network communication system according to claim 18, wherein the plurality of different types of traffic comprises at least two of network traffic, storage traffic, interprocess communication (IPC) traffic and cluster traffic.

33. A method for communication, the method comprising:
in a data center:

accessing a storage system over a single fabric, wherein said single fabric comprises a single layer 2 (L2) connector coupled to a single integrated convergent network controller chip that is enabled to concurrently process a plurality of different types of traffic; and

accessing one or more of a cluster and a network over said single fabric.

34. The method according to claim 33, wherein said accessing of said storage system, over said single fabric are performed over a single Ethernet connector of a server in the data center.

35. The method according to claim 33, wherein said single integrated convergent network controller chip coupled to the single Ethernet connector has a single Internet protocol (IP) address.

36. A system for communication, the system comprising:

a single integrated convergent network controller chip that enables concurrent hardware, firmware and software processing functionalities of a plurality of different types of traffic that are received via a single layer 2 (L2) connector that is communicatively coupled to a plurality of servers via a single fabric.

37. The system of claim 36, wherein said single integrated convergent network controller chip comprises a layer 2 network interface card (L2 NIC), a

transmission control protocol (TCP) processor, an iSCSI processor, a remote direct memory access (RDMA) processor and a Management Agent processor.

38. The system of claim 36, wherein said plurality of different types of network traffic comprises at least two of a network traffic, storage traffic, interprocess communication (IPC) traffic and cluster traffic.

39. A method for communication, the method comprising:
concurrently providing, via a single integrated convergent network controller chip, hardware, firmware and software processing functionalities of a plurality of different types of traffic that are received via a single layer 2 (L2) connector that is communicatively coupled to a plurality of servers via a single fabric.

40. The method of claim 39, wherein said single integrated convergent network controller chip comprises a layer 2 network interface card (L2 NIC), a transmission control protocol (TCP) processor, an iSCSI processor, a remote direct memory access (RDMA) processor and a Management Agent processor.

41. The method of claim 39, wherein said plurality of different types of network traffic comprises at least two of a network traffic, storage traffic, interprocess communication (IPC) traffic and cluster traffic.

EVIDENCE APPENDIX
(37 C.F.R. § 41.37(c)(1)(ix))

- (1) United States Pub. No. 2001/0037406 ("Philbrick"), entered into record by the Examiner in the 10/6/2009 Final Office Action.
- (2) United States Pub. No. 2002/0059451 ("Haviv"), entered into record by the Examiner in the 10/6/2009 Final Office Action.
- (3) Microsoft Winsock Direct and Protocol Offload on SANs, 03/03/2001 ("Microsoft"), entered into record by the Examiner in the 10/6/2009 Final Office Action.

RELATED PROCEEDINGS APPENDIX
(37 C.F.R. § 41.37(c)(1)(x))

The Appellant is unaware of any related appeals or interferences.